PRECISE AND EFFICIENT FORMATION KEEPING AT EARTH-SUN L2 FOR STARSHADE MISSIONS

Thibault L. B. Flinois,[*] Daniel P. Scharf,[†] Carl R. Seubert,[†] Michael Bottom,[‡] and Stefan R. Martin[‡]

Current starshade concepts for imaging exo-Earths would operate at the Earth-Sun L2 point and consist of a starshade flying in formation tens to hundreds of thousands of kilometers from its telescope. The starshade would need to be aligned to the meter level across the line of sight from telescope to target star. This paper reports work aimed at maturing the technology readiness level of starshade formation sensing and control for the starshade science phase, during which precision alignment is required. First, a formation control algorithm that provides high efficiency for science observations is presented. Then, other elements of an end-to-end starshade GNC system for the precision alignment phase are outlined: the enabling formation fine lateral sensor using pupil-plane images, an RF communication link and ranging sensor, the formation estimator, the longitudinal formation control algorithm, and a thrust allocation algorithm. A high-fidelity simulation environment, which includes sensor, actuator, and dynamics models, is also summarized. Finally, for a representative observation scenario with WFIRST, this GNC system for precision alignment is demonstrated in a Monte Carlo simulation, robustly achieving observationally efficient meter-level control.

INTRODUCTION

Starshades – precisely shaped external occulters flying in formation with a telescope – provide a method for direct observation and spectroscopic characterization of Earth-like exoplanets.[1,2,3,4] Historical perspectives are included in References 3 and 5. In the 2000s, flagship-class starshade missions were proposed.[6,7] In this decade, NASA is studying both smaller, probe-class missions with 20 to 40 m diameter starshades that could rendezvous and operate with the Wide Field Infrared Survey Telescope (WFIRST)[1,8] and a larger, approximately 80 m diameter flagship-class mission for the 2020 Decadal Survey.[9] The operational concept expected for these missions is summarized in Figure 1. To support the objectives of these missions, NASA's Exoplanet Exploration Program is maturing key starshade technologies to Technology Readiness Level 5 (TRL 5) and has combined many of these maturation efforts into the S5 activity ("Starshade to TRL 5").[10] The work reported here was done as part of S5 and is intended to demonstrate robust, meter-level, lateral formation control during science observations of exoplanets ("Science phase" in Figure 1). This level of control is required to maintain a starshade (and its high-contrast shadow) aligned with the

telescope. The other phases of starshade formation flying, such as initialization, acquisition, and retargeting, while requiring significant engineering, are not considered technology risks as the performance requirements are much looser and the corresponding formation sensors would be more standard (e.g., a camera measuring bearing to a beacon).[11]
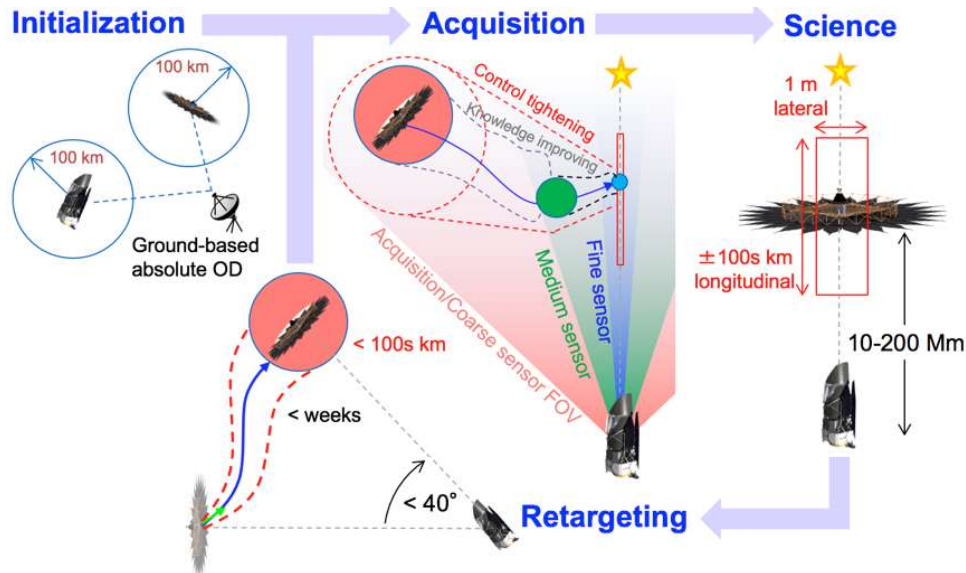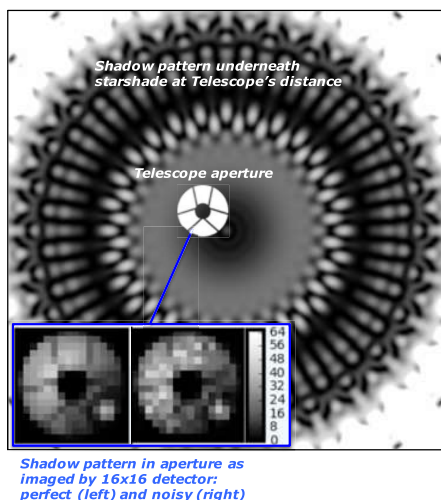

Figure 1. Starshade operational phases.

As discussed in Reference 11 and noted elsewhere,[12,13] the micro-g-level differential acceleration between a starshade and telescope separated by ~100 Mm at the Earth-Sun L2 point is comparable to that experienced by spacecraft docking in LEO, which regularly achieve 10 cm-level relative control. Hence, the principal technological challenge for formation control in the observational "Science phase" is developing a so-called fine lateral sensor that can measure sub-meter offsets of a starshade relative to an occulted target star at tens to hundreds of megameter separations. Many fine lateral sensing approaches have been studied to achieve this goal; References 11 and 14 summarize several. In particular, References 12 and 15 leverage the residual Arago spot in the pupil plane: either centroiding the image or fitting the image to an analytic approximation of the Arago spot. Recently, Reference 16 developed a variation on this approach, fitting small, 16x16-pixel images to a pre-computed library of pupil plane images. This approach is sufficiently precise (better than 10 cm 3σ within 1 meter of alignment and better than 25 cm 3σ across the entire shadow for an 8[th] magnitude star with 1 s exposures), is simple (based on least squares), functions without modification from zero lateral offset to the radius of the shadow, and functions with obscuration by a secondary. The approach and its performance have also been validated in a scaled hardware testbed[17]. This fine lateral sensor is the foundation of the GNC system for the Science phase, and more details are provided in the next section.

Previous work on closed-loop control for the Science phase has focused on feasibility studies with lower fidelity simulations, lower fidelity GNC systems, or continuous actuation (e.g., using electric propulsion). Electric propulsion on a starshade is likely not suitable for the Science phase as light emitted by electric thrusters or sunlight reflected off of their plumes could corrupt observations. As plumes from chemical thrusters could also interrupt observations, chemical thruster firings should be as infrequent as possible. Regarding previous work, Reference 11 uses a more complicated focal-plane fine lateral sensor, an extended Kalman Filter (EKF), and a single-axis deadbanding algorithm in a lower-fidelity, single-axis simulation. Reference 12 also

has a closed-loop single-axis simulation with continuous actuation and with an estimator replaced by sensor noise. Reference 13 includes an open-loop sensitivity study of single-axis deadbanding with various differential accelerations, thruster sizes, and firing times, but without sensor noise. Reference 18 reports 3-axis simulations using LQR, an EKF, and continuous actuation.

The contributions of this paper are 1) to present flight-like algorithms for the Science phase of starshade operations and 2) to demonstrate that these algorithms robustly achieve the required lateral control performance in a high-fidelity simulation environment. The paper is structured as follows. First, the fine lateral sensor and its testbed-validated error model are described. Second, a new, efficient, and robust two-axis deadbanding control algorithm is presented. Third, the remaining formation GNC algorithms are outlined, including the longitudinal control algorithm, the formation estimator that accounts for variably delayed measurements communicated from the telescope to the starshade, and the thrust allocator. Fourth, the dynamics, sensor, and actuator models of the high-fidelity simulation environment are described. Fifth, the results of Monte Carlo simulations are presented showing that both the required lateral control performance and effectively optimal drift time between thruster firings are achieved. Finally, conclusions of the work are given.

OVERVIEW OF FINE LATERAL SENSOR

Shadow pattern underneath starshade at Telescope's distance

Telescope aperture

Shadow pattern in aperture as imaged by 16x16 detector: perfect (left) and noisy (right)

Figure 2. Pupil-plane sensing approach: pupil plane images of a small portion of the out-of-band diffraction pattern are captured and matched to a precomputed library to compute lateral relative position.

The GNC framework presented here relies on a pupil-plane sensor that images the residual starlight diffracted by the starshade in a pupil imager inside the telescope. The concept of sensing starshade displacement via leaked starlight has long been discussed in the literature[12,14,15,19]. A starshade converts the flat incident wavefront into a complex pattern at the pupil of the telescope, which consists of a weak central Arago spot tens of centimeters wide and a set of radial structures that are due to the starshade petals (Figure 2). This pattern is extremely faint ($\sim 10^{-10}$ flux ratio) at the science wavelength, but brightens substantially outside the science bands ($\sim 10^{-3}$ flux ratio).

In Reference 16, a library of ideal images of the optical field behind a representative starshade is first generated for a grid of lateral offsets. These reference images have their average brightness subtracted out to ensure image matching can focus on features and not image brightness. Second, for a separate grid of lateral offsets, example pupil-plane images are generated with optical path effects, pixelation, and detector noise. These images also have their (noisy) average brightness subtracted out. Finally, for each noisy image, least-squares is used to find the closest image match in the library, thereby giving the corresponding lateral offset (Figure 2). This lateral sensing approach was validated analytically, numerically, and experimentally in Reference 17. The performance of this sensor varies with the radial position of the starshade with respect to the telescope-star axis (see spatially varying sensor performance in left part of Figure 3) and in general the covariance matrix is not diagonal (i.e. it has tilted ellipses).

This sensor model was used in the present work. For added conservatism, the sensor noise predicted by Reference 17 was scaled such that for all lateral offsets, no measurement has a standard deviation smaller than 10cm (i.e. 30 cm, 3σ) in any direction (right part of Figure 3). This approach

maintains the spatial variation of the sensor accuracy as a function of the Starshade lateral offset, but with extremely conservative performance.
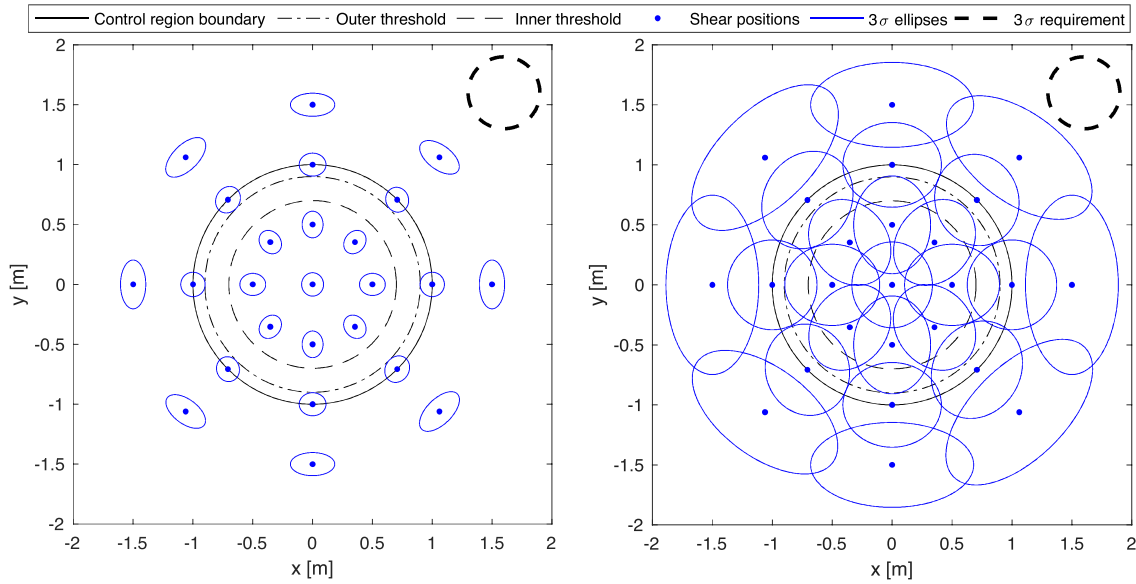
Figure 3. Lateral sensor model: Left: original testbed-validated model. Right: Scaled, extremely conservative model used in formation flying simulations (in target frame: z axis aligned with telescope-to-target-star axis). Sensing 3σ requirement of 30cm shown in top-right corner. Control requirement and (subsequently described) thresholds for deadbanding control algorithm also shown.

FORMATION KEEPING ALGORITHM

In this section, the control algorithm for the Science phase is presented. The overall control objective is to maximize the drift time between thruster firings while remaining inside the required circle. For the timescales in question, the disturbing relative lateral acceleration vector can be taken as constant. Simply put, the goal is to optimally deadband within a circle. It appears that this problem has not been previously addressed: typical deadbanding for spacecraft attitude or docking/berthing control is per-axis, where per-axis burns are not synchronized. For starshades, unsynchronized burns would further interrupt science observations and would result in a rectangular deadband region, which does not conform well to the circular geometry of the shadow.

The sections below introduce the trajectory that gives the maximum drift time in the circle and the Δv that gives this trajectory. A subtlety for circular deadbanding is that the maximum achievable drift time depends on the initial position on the circle. The globally optimal trajectory is entirely along the diameter of the circle that is aligned with the acceleration vector. Given an initial position on the circle, the deadbanding Δv therefore maximizes drift time, while targeting the starting point of the globally optimal drift trajectory. As a result, the algorithm seamlessly produces the globally optimal drift: no nonlinear switching logic is needed. However, two thresholds (i.e., smaller circles) are introduced so that the starshade remains in the required circle in the presence of sensor noise and Δv execution errors.

Deadbanding trajectory

The algorithm presented here assumes the lateral components of the relative acceleration, velocity, and position vectors are being estimated on board. This allows defining the acceleration-aligned lateral coordinate frame shown in Figure 4, where the trajectory, immediately following a burn, starts at time $t = 0$. The constant relative lateral acceleration vector is $[0, -a]^T$, where $a > 0$.

4

(In practice, when the estimated acceleration is below a certain threshold, a small yet accurately achievable Δv can simply be applied towards the center of the circle.) For clarity, dynamics are nondimensionalized by scaling lengths by the deadband radius R_d, velocities by $\sqrt{aR_d}$, and time by $\sqrt{R_d/a}$. The non-dimensional velocity and position vectors are thus given by $[u(t), v(t)]^T = [u_0, -t + v_0]^T$ and $[x(t), y(t)]^T = [u_0 t + x_0, -t^2/2 + v_0 t + y_0]^T$ respectively, where $[x_0, y_0]^T$ and $[u_0, v_0]^T$ are the initial position and velocity vectors respectively. It can be shown that the trajectory that leads to the longest possible drift time within the deadband region is, of course, entirely along the diameter of the circle that is aligned with the acceleration vector. This globally optimal trajectory has its start and end positions located at "the well" $[0, -1]^T$ and its peak at $[0, 1]^T$. The associated nondimensionalized drift time is $t_{f,\max} = 4$.
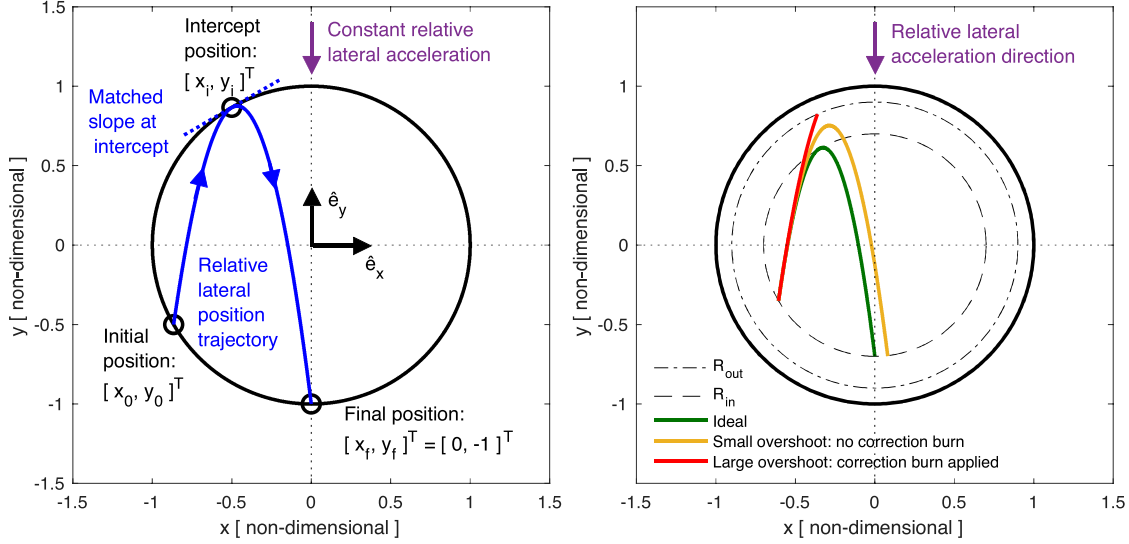
Figure 4. Left: Control approach overview. Right: Double threshold approach. (Both in acceleration-aligned target frame: z axis aligned telescope-to-target-star axis, negative y axis aligned with lateral component of relative acceleration vector.)

In general, however, the initial position $[x_0, y_0]^T$ is not located at the well. It can also be shown numerically that the trajectory that "locally" maximizes drift time – i.e. maximizes drift time given a starting position on the circle – does not end at the well. Locally maximizing drift time would therefore lead to several globally suboptimal trajectories, even with no errors. Instead, the chosen approach is to determine the trajectory that, given $[x_0, y_0]^T$, maximizes drift time, *while also always targeting the well*. This means that although the drift time after the first burn may be locally slightly suboptimal (always by less than 10% compared to local optimal), all the following ones will in theory be globally optimal. In practice, this allows the algorithm to seamlessly produce trajectories that are effectively globally optimal even in the presence of execution error, without the need for any nonlinear switching logic. Given a starting position on the deadband boundary and a final position at the well, the trajectory with the longest drift time can be shown to intercept the deadband boundary at a third point and is tangent to it at that point (see Figure 4, left).

Deadbanding Thruster Burn

The desired trajectory satisfies the following conditions: the position of the intercept is on the deadband boundary and on the parabolic trajectory, the intercept is tangent to the deadband boundary, and the final position (at the well) is also on the parabolic trajectory. These conditions

lead to a quartic equation, which was solved analytically for the four unknowns $[x_i, y_i, u_0, v_0]$. For $x_0 < 0$, the intercept position was found to be:

$$[x_i, y_i]^T = \begin{cases} [x_0, y_0]^T, & \text{for } y_0 > \frac{1}{2}, \\ \frac{1}{\sqrt{2}} \left[-\sqrt{1+y_0}, \quad \sqrt{1-y_0} \right]^T, & \text{for } y_0 < \frac{1}{2}, \end{cases} \tag{1}$$

and the initial velocity is given by:

$$[u_0, v_0]^T = \left[\frac{1}{\sqrt{2}} \sqrt{y_i (1 - y_i)}, \quad v_i + t_i \right]^T, \tag{2}$$

where $v_i = -u_0 x_i / y_i$ is the vertical velocity at the intercept and $t_i = (x_i - x_0) / u_0$ is the time required to reach the intercept. For the $y_0 > 1/2$ case, note that $t_i = 0$, i.e. the initial position is the intercept. For $x_0 > 0$, the sign of x_i for the $y_0 < 1/2$ case and the sign of u_0 need to be inverted.

Double Threshold Approach

In practice, estimation and thruster errors often lead to slightly suboptimal yet acceptable trajectories. To ensure that thruster firings are not triggered unnecessarily (to avoid interrupting science) but that the trajectory always remains within the deadband boundary, an inner and an outer threshold are defined based on the values of the estimated relative lateral position $\hat{\mathbf{r}}$, velocity $\hat{\mathbf{v}}$, and acceleration $\hat{\mathbf{a}}$ vectors. A burn is triggered if the relative lateral position drifts outside the outer threshold with radius R_{out}, while getting further from the origin (so no burn is triggered while re-entering). $R_{\text{out}} < R_d$ is required to give the control system sufficient time to trigger the burn before R_d is crossed. In addition, a burn is triggered if the relative lateral position drifts outside the inner threshold with radius R_{in}, while getting further from the origin, but only in the lower half (negative y) of the circle. This allows overshoots beyond R_{in} in the top (positive y) half of the circle as long as these remain within R_{out} (see Figure 4, right). Formally, a burn is triggered if the following criteria are met:

- Outer threshold: $\quad \|\hat{\mathbf{r}}\| > R_{\text{out}}, \quad$ and $\quad \hat{\mathbf{r}} \cdot \hat{\mathbf{v}} > 0$.

- Inner threshold: $\quad \|\hat{\mathbf{r}}\| > R_{\text{in}}, \quad$ and $\quad \hat{\mathbf{r}} \cdot \hat{\mathbf{v}} > 0, \quad$ and $\quad \hat{\mathbf{r}} \cdot \hat{\mathbf{a}} > 0$.

The thresholds were tuned to $R_{\text{in}} = 0.7 R_d$ and $R_{\text{out}} = 0.9 R_d$.

OVERVIEW OF REMAINING GNC ALGORITHMS

Longitudinal Control

Given the low relative acceleration ($10\ \mu\text{m/s}^2$ level) compared to the longitudinal control region size (hundreds of kilometers, see Figure 1), the longitudinal control boundary is likely not to be crossed for hours or days even without longitudinal control. However, when needed, execution errors from comparatively large longitudinal burns can significantly disturb lateral drift trajectories. Therefore, a longitudinal control algorithm was developed that increases longitudinal drift time but that has negligible impact on lateral control: synchronized longitudinal rate damping. When the relative longitudinal speed is above a threshold, a longitudinal Δv is applied with the next lateral control Δv. The longitudinal Δv magnitude is also limited to another threshold.

Estimation

A Kalman filter was implemented to estimate the three-dimensional relative position, velocity, and acceleration between the two spacecraft. The Kalman filter is fed only with delayed and noisy sensor measurements and their associated time-tags. Since the rate-of-change of the relative

acceleration is slow compared to the drift time between burns, the filter dynamics use a constant, three-dimensional acceleration.

The process noise matrix was tuned to ensure the response of the relative acceleration estimation was sufficiently fast to track the slow evolution of the relative acceleration vector as the two spacecraft move along their orbit, but sufficiently slow to reject higher frequency disturbances due to sensor noise and starshade-spin-induced solar radiation pressure (SRP) fluctuations. Further, Δv's were modelled as process uncertainties. Thus, when firing, a different process noise matrix was used. The measurement noise matrix was designed to account for longitudinal sensor precision (RF ranging with 500 m 3σ), time-tagging uncertainties, clock drift uncertainty, and spatially-dependent lateral sensor precision (see Figure 3).

Longitudinal and lateral position measurements become available to the estimator in a delayed and asynchronous manner. Several exact and approximate methods for dealing with such delays exist[20,21,22,23]. Here, an essentially brute-force Kalman filtering scheme was implemented that ensures that all delayed measurements (up to several seconds) are processed optimally. Past state estimates and covariance matrices are stored and used to recompute the state history when delayed measurements become available. This approach is believed reasonable for now as the starshade control update rate is of the order of 1 Hz or less and the filter state is not large.

Attitude knowledge is required for thrust allocation. Since this knowledge is not driving, raw star-tracker measurements are simply propagated using gyro measurements to the desired time.

Thrust Allocation

Thrust allocation is non-trivial as the thrusters are not aligned with body axes and the representative starshade is spin-stabilized at 1/3 rpm. The QPCAP (Quadratic Programming algorithm for the Control Allocation Problem)[24] thrust allocation algorithm was used to realize torque-free Δv's with optimal thruster on-times. QPCAP flew on the Mars Science Laboratory (MSL) descent stage and has been used extensively in formation flying studies and demonstrations[25].

HIGH-FIDELITY SIMULATIONS

The algorithms presented above were tested in a high-fidelity simulation environment described in this section. This MATLAB-based environment propagates the position of a telescope and a starshade spacecraft from initial conditions in three-dimensional space, under the influence of SRP and the gravity from Earth, Sun, Moon, and solar system planets. The demonstration scenario is a representative starshade rendezvous mission with WFIRST[10].

Simulation Dynamics

The representative starshade is composed of a 10 m diameter disk with 24 petals of 10.66 m^2 area each, distributed around the disk circumference. The starshade spacecraft (including the bus and the starshade) dry mass is 1200 kg. For conservatism, only 100 kg of propellant is on board so as to increase the influence of mass uncertainty. A 2% propellant mass knowledge uncertainty is standard for typical spacecraft. Based on an assumption of 2500 kg maximum propellant mass, 50 kg (3σ) mass uncertainty was assumed.

The orbital propagation uses a constant time-step, 5[th] order Runge-Kutta ODE solver. The ephemerides of the bodies were obtained from JPL's SPICE library[*] and validated against MONTE simulations (JPL's high-fidelity mission design tool, used on several NASA deep space missions[†]).

A standard SRP force model was implemented with absorptive (92%), specular reflective (4%), and diffusive reflective (4%) components. SRP acceleration is considered to be negligible for WFIRST due to the comparatively small surface area.

Attitude control is modeled only to the extent that it affects the formation keeping. Specifically, the starshade is spin-stabilized (constant spin rate of 2°/s) and when the starshade axis is not aligned with the longitudinal axis, thrust allocation is stressed, contributing to Δv execution errors. The required alignment of the starshade symmetry axis and the formation longitudinal axis is 1°. Therefore, the starshade spin axis is prescribed to be offset at a constant angle of 1°. (Note that attitude errors of this magnitude do not have a significant effect on the lateral sensing performance.) This was implemented by propagating the torque-free attitude of the starshade[‡] from initial conditions chosen to lead to nutation-free motion. This leads to the spin axis precessing at a constant rate of about 4°/s around the starshade-to-telescope axis.
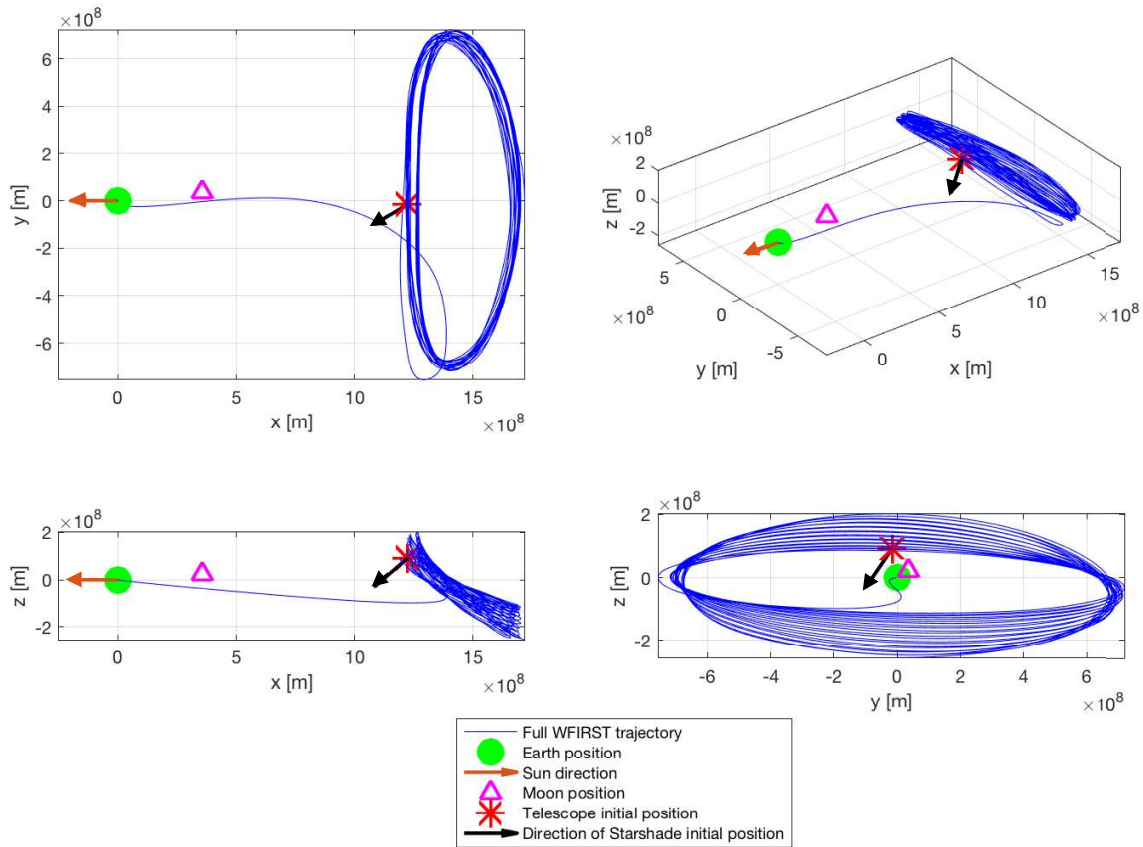
Figure 5. Worst-case initial formation geometry in Earth-centered rotating frame: x axis aligned with Sun-Earth vector and y axis in the plane of the Earth velocity vector. Note that only the telescope-starshade vector *direction* is shown as the Science phase separation is too small to show to scale. Representative telescope trajectory courtesy of C. Webster and WFIRST mission design team.

The simulated formation geometry was selected to maximize the relative lateral acceleration. A representative WFIRST trajectory was provided by NASA Goddard Spaceflight Center. This trajectory along with relative directions of the starshade were searched numerically. Intuitively, the worst-case geometry consists of the telescope closer to the Earth and Moon and the starshade at its maximum operating range of ~40 Mm (for a WFIRST observation scenario). The worst Earth-telescope-starshade angle is close to 45°: the relative dynamics are driven by Earth gravity and this angle strikes a balance between having the telescope, starshade and Earth aligned (0°) and having the starshade and telescope equidistant to Earth (90°), both of which lead to zero lateral relative acceleration. The numerically identified worst-case formation geometry is illustrated in Figure 5. The corresponding relative lateral acceleration due to SRP and the gravity of Earth, Sun, Moon, and solar system planets is 15.2 $\mu m/s^2$ (1.55 μg), for a representative observation scenario with WFIRST.

Thruster Model and Configuration
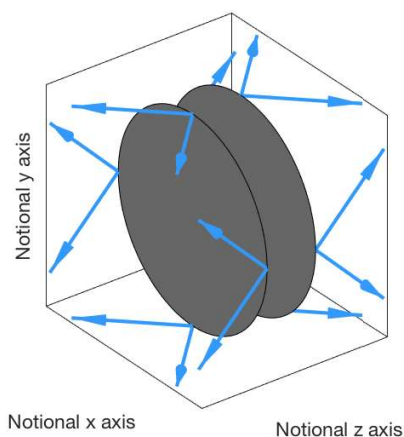
Figure 6. Schematic views of thruster configuration, shown by thruster plume directions, (in starshade body-fixed frame: x axis aligned with starshade spin axis).

A 16-thruster configuration was chosen for the starshade, as illustrated in Figure 6, based on balancing the needs of fine formation control and applying retargeting Δv. The thrusters are all nominally oriented at a 45° angle to the starshade spin axis and their projection onto the plane normal to the spin axis is either oriented in the ±x or ±y directions. Thrusters used in the simulations are similar in performance to the flight-qualified Aerojet Rocketdyne R-6D bipropellant thrusters, considered to be representative. These thrusters have a nominal thrust of 22 N, which enables short burns relative to the starshade spin rate.

Thruster errors were implemented as follows: known thrust direction and magnitude biases of 1° (3σ) and 1% (3σ); unknown biases of 0.75° (3σ) and 2% (3σ); and unknown "pulse-to-pulse" errors of 0.1° (3σ) and 1% (3σ). Each thruster was assumed to have a minimum on-time of 5 ms and a quantization of 0.5 ms. A conservative, one second delay was applied between any command and the start of the corresponding burn.

Sensor Models

The lateral sensor model is described in a previous section. The longitudinal sensor uses the S-band inter-spacecraft communication system and is assumed to have 500 m 3σ precision. Starshade attitude sensing was implemented with a 20 arcsecond 3σ attitude error and a 0.01°/s 3σ angular rate error. This is consistent with standard performance from flight-proven attitude sensors. For both longitudinal and lateral sensors, conservative clock drift, measurement time variations, time-tagging errors, and measurement delays were also included.

Monte Carlo Set Up

A set of Monte Carlo simulations was run using the previously presented uncertainties and random initial relative positions and velocities, to demonstrate the robustness and statistically assess the performance of this starshade GNC system. The main performance parameter is the drift time between Δv's, given that the 1 m radial formation-flying control requirement is met. Each simulation was 6 hours long, allowing time for ~25 lateral control Δv's, with steady-state

performance beginning after ~3 Δv's as measured by average drift time. A total of 60 simulations was run. At 60 runs, the resulting distribution of drift times was found to be approximately normally distributed up to about two standard deviations.

SIMULATION RESULTS

In all cases, the lateral offset of the starshade remained within the required circle. Figure 7 shows a typical deadbanding trajectory over one hour, starting from a random initial position and relative velocity. As expected, the lateral trajectory remains inside the inner threshold, and from the second burn onwards, the trajectory remains close to the optimal trajectory, i.e. almost aligned with the relative lateral acceleration.
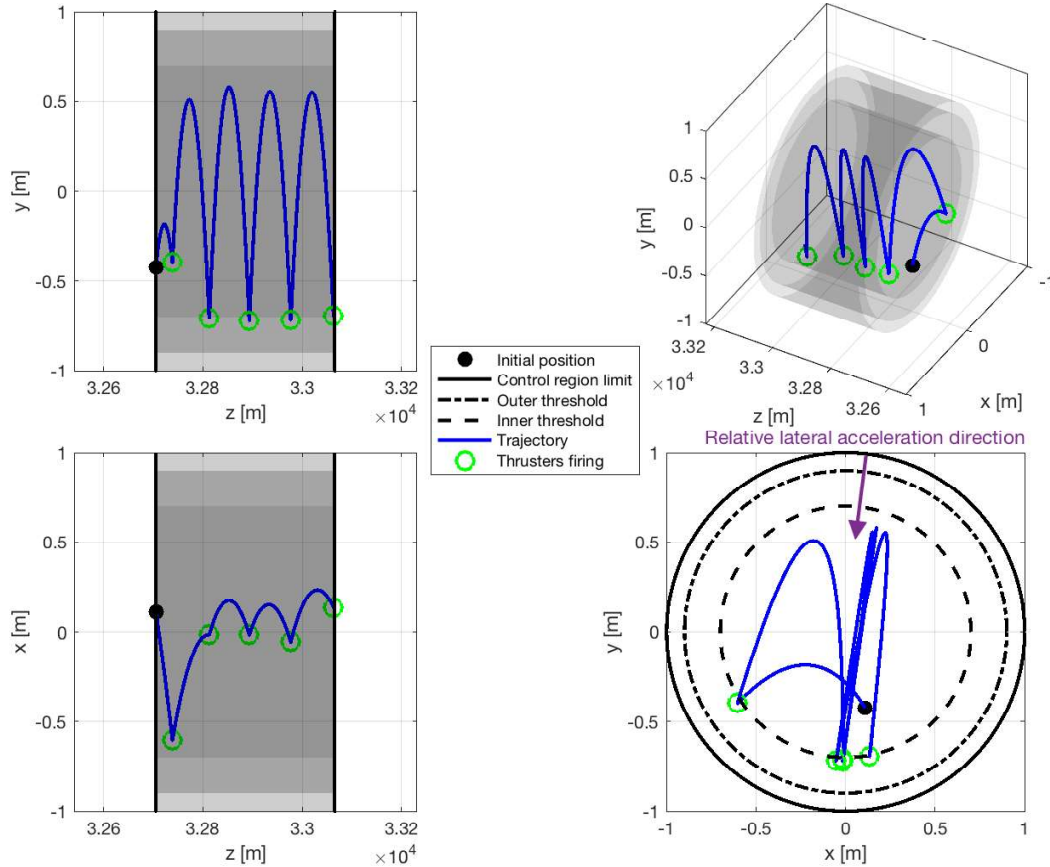

Figure 7. Example of a typical three-dimensional trajectory over one hour, showing the two thresholds and the control region boundary (in target frame – not acceleration aligned).

Figure 8 shows the statistics of the resulting drift times across all 60 runs starting at the third Δv, compared to the drift time expected for the ideal deadbanding trajectory. The mean drift time across all simulations was found to be 850 s. This value is about 1% lower than the drift time expected for ideal deadbanding about R_{in}. This bias is likely due to the fact that all trajectories with any directional error in Δv or with starting positions anywhere else than exactly at the well have reduced drift times.

Figure 8 also shows that choosing $R_{\mathrm{in}} = 0.7 R_d$ reduces the ideal drift time by about 16%. On the other hand, in rare cases (<1% of the burns), the Δv execution error is large enough for the outer threshold to be reached near the "top" of the deadband (see red curve in Figure 4, right), thereby

10

triggering a corrective burn and approximately halving the drift time. This behavior was found to take place when the estimate of the starshade's mass is much higher than the actual mass, leading to force commands that are systematically much larger than the desired Δv's. In the present set of simulations, this "two-sided" deadbanding occurred several times for exactly one of the simulations that had a mass estimation error of 41 kg, which is ~2.5σ or ~3% of wet mass. It can be shown analytically that mass uncertainty is the driving error source for overshoot. Hence, there is a trade-off: increasing R_{in} and R_{out} increases the nominal drift time but also increases the probability of double-sided firing. In practice, ground calibration – or even using the overshoot itself during Science as a calibration method – should to reduce the wet mass uncertainty to much better than 3%.
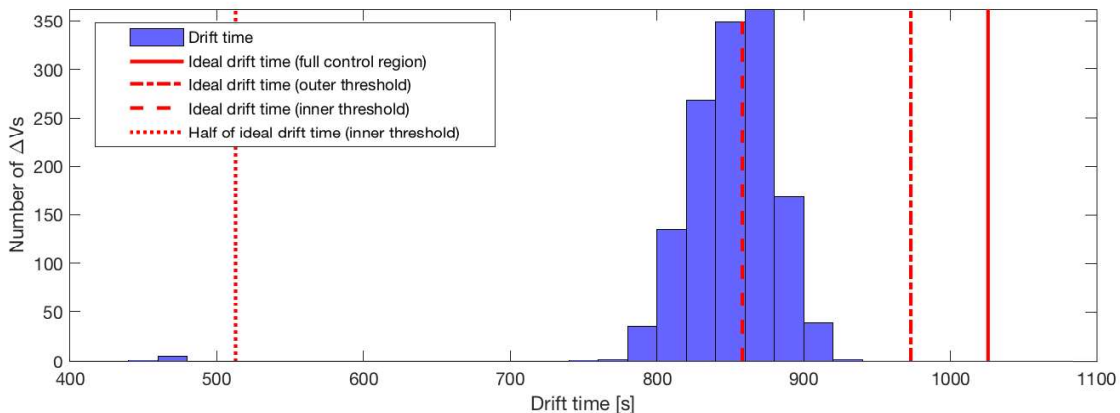
Figure 8. Histogram of drift times resulting from the Monte Carlo simulations.

CONCLUSIONS

A GNC framework for fine formation keeping between a starshade and a space telescope in halo orbits around the Earth-Sun L2 Lagrange point was introduced and demonstrated in a high-fidelity simulation environment. The new elements of this framework are a pupil-plane image-matching fine lateral sensor and a circular deadbanding algorithm that maintains the starshade within a 1 m circle from the telescope line of sight. Monte Carlo simulations with conservative assumptions and with the most stressing lateral acceleration for a representative WFIRST starshade rendezvous scenario confirmed the robustness of this GNC framework and showed that it leads to essentially optimal drift times between Δv's (given the control tuning and the relative acceleration), thereby maximizing scientific observational efficiency.

Given the fidelity of the simulation environment – including the testbed-validated sensor noise model – and GNC algorithms and the conservatism in both the lateral sensor performance, sensor delays, and thruster errors, it is believed that the simulation results demonstrate formation control for the Science phase to TRL5. A more detailed report on this work is being reviewed by a NASA independent review board and will be made publicly available online.[*] The retargeting phase has also been analyzed in detail and a publication is pending. Future work will also look at simulations of the Acquisition phase.

---

[*] https://exoplanets.nasa.gov/exep/technology/starshade/

ACKNOWLEDGMENTS

[1] S. Seager et al., "Imaging and Spectra of Exoplanets Orbiting our Nearest Sun-Like Star Neighbors with a Starshade in the 2020s," *White Paper submitted to National Academy of Sciences*, 2018.

[2] M. Levine and R. Soummer, "Overview of Technologies for Direct Optical Imaging of Exoplanets," *ASTRO2010 Community White Paper*, 2009.

[3] W. Cash et al., "External Occulters for Direct Observation of Exoplanets: An Overview," *Proc. SPIE*, Vol. 6687, 2007.

[4] N.J. Kasdin et al., "Imaging Extrasolar Planets via Formation Flying with an External Occulter," *Adv. Astro. Sci.,* Vol. 132, pp. 115-127, 2008.

[5] M. Kochte et al., "External Occulter Operations Requirements," *AIAA Space OPS Conf.*, 2004.

[6] W. Cash et al., "The New Worlds Observer: The Astrophysics Strategic Mission Concept Study," *Proc. SPIE*, Vol. 7436, 2009.

[7] N.J. Kasdin et al., "Occulter Design for THEIA," *Proc. of SPIE Vol. 7440*, 2009.

[8] S. Seager et al., "Exo-S: Starshade Probe-Class Exoplanet Direct Imaging Mission Concept, Final Report," 2015.

[9] B.S. Gaudi et al., "The Habitable Exoplanet Observatory (HabEx)," *Proc. SPIE*, Vol. 10698, 2018.

[10] P. Willems et al., "Starshade to TRL5 (S5) Technology Development Plan," NASA Report, 2018. url: https://exoplanets.nasa.gov/internal_resources/1033

[11] D.P. Scharf et al., "Precision Formation Flying at Megameter Separations for Exoplanet Characterization," *Acta Astronautica*, Vol. 123, 2016, pp. 420–434.

[12] M.C. Noecker, "Alignment Control for an External-Occulter Terrestrial Planet Finder Mission," *Adv. Astro. Sci.*, Vol. 137, pp. 203-220, 2010.

[13] D. Sirbu, C.V. Karsten, J.N. Kasdin, "Dynamical Performance for Science-Mode Stationkeeping with an External Occulter," *Proc. SPIE*, Vol. 7731, 2010.

[14] M.C. Noecker, "Alignment of Terrestrial Planet Finder Starshade at 20-100 Megameters," *Proc. SPIE*, Vol. 6693, 2007.

[15] A. Harness, "Enabling Formation Flying of Starshades for the Search of Earth-Like Exoplanets," *8th Intl. Workshop on Spacecraft Constellations and Formation Flying*, 2015.

[16] M. Bottom et al., "Precise Starshade Stationkeeping and Pointing with a Zernike Wavefront Sensor," *Proc. SPIE*, Vol. 10400, 2017.

[17] M. Bottom et al., "S5: Starshade Technology to TRL5 - Milestone 4 Final Report", under review by NASA independent review board, to be released online: https://exoplanets.nasa.gov/exep/technology/starshade/

[18] D. Sirbu et al., "Stationkeeping for an Occulter-Based Exoplanetary Imaging Mission," *Adv. Astro. Sci.*, Vol. 137, pp. 129-148, 2010.

[19] W. Cash, "The New Worlds Observer: the Astrophysics Strategic Mission Concept Study," *EPJ Web Conf.,* Vol. 16, 2011.

[20] H.L. Alexander, "State estimation for distributed systems with sensing delay" *Proc. SPIE*, Vol. 1470, 1991.

[21] T.D. Larsen et al., "Incorporation of Time Delayed Measurements in a Discrete-time Kalman Filter," *Proc. IEEE Conf. Decis. Cont.*, pp. 3972-3977, 1998.

[22] S. Challa, R.J. Evans, X. Wang, "A Bayesian solution and its ap- proximations to out-of-sequence measurement problems," *Inform. Fusion*, Vol. 4, pp. 185-199, 2003.

[23] Y. Bar-Shalom, H. Chen, M. Mallick, "One-Step Solution for the Multistep Out-of-Sequence-Measurement Problem in Tracking," *IEEE Trans. Aerosp. Electron. Syst.*, Vol. 40, No. 1, pp. 27-37, 2004.

[24] G. Singh, "QPCAP: A Quadratic Programming Algorithm for the Control Allocation Problem," *JPL Engineering Memorandum 3457-03-01*, 2003.

[25] D.P. Scharf, J.A. Keim, F.Y. Hadaegh, "Flight-like Ground Demonstrations of Precision Maneuvers for Spacecraft Formations - Part II," *IEEE Syst. J.*, Vol. 4, No. 1, 2010, pp. 96–106.